

Second Declaration of George T. Ligler

Exhibit 9

19. NOV. 2004 7:41

44_01937_5466492

NO. 6111 P. 2


Vol. 12, No. 7
Dec. 13, 1982
ISSN: 0364-5517

The Seybold Report

on Publishing Systems

The Hell Chromacom:

A Tool for Today, a Vision for Tomorrow



HELL IS THE LARGEST SUPPLIER of color-separation scanners. (Hell's own estimates are that it has 60%-65% of the U.S. market, and 52% of the market worldwide.) For many quality-conscious color printers these machines have become the preferred means of producing sized and screened color separations from original transparencies or color prints. In the last decade, Hell, Crosfield, and other companies began work on digital color systems which would allow manipulation and assembly of color images to be performed in between the input scanning and output writing operations of a color separation scanner.

Actually, as often happens when new technology hits an industry, the key innovator in this field has been neither of the two established firms but Scitex, an "upstart" which has entered the graphic arts industry from other fields and has brought along its own technology and insights. Both Crosfield and Hell were hard at work on development of digital color systems long before Scitex appeared on the scene. But both have clearly been influenced by what Scitex has done and by the way in which Scitex has been able to capture the imagination of the marketplace.

Scitex brought its system to market first, Crosfield, whose highly modular approach made it possible to install systems for page layout and assembly without any color preview or correction facilities, followed quickly. However, Crosfield has had difficulty getting the final pieces of its system to the field.


Hell, by contrast, chose to tackle a full-function system, rather than to proceed modularly. It has been building a substantial base of installations over the last two years.

In certain areas (particularly real-time image sizing and rotation and manual switching of disks) the Hell system is still somewhat less sophisticated than Scitex's. However, the system is now selling very well indeed and sales momentum is encouraged by virtue of the large installed base of Hell scanner users. The current high value of the dollar in relation to the Deutschmark is also a positive factor. (Scitex prices are based on U.S. dollars so that exchange rates, at least in relation to the mark, are not relevant.) But beyond this, the Hell system appears to be a sensible and practical production tool.

For the future, Hell, like Scitex and Crosfield, intends to incorporate the ability to generate and output text as well as graphics. And, like Scitex, it intends to move "upstream" in the production cycle with development of a less-expensive workstation which can be used for design and page layout.


Electronic color systems have caused the blurring of the traditional craft distinctions between color separation, retouching, and stripping. In the same way, new developments in incorporating text are blurring the distinction between color operations and typesetting. The design and plate-making functions are next on the list of areas to be incorporated. We are beginning to see the emergence of total systems which will handle all pre-press functions in an electronic environment.

In our coverage of the Print '80 show at which these color systems burst onto the U.S. scene (Vol. 9, No. 16/17), we noted that they foreshadowed "the day when handling of text and graphics are far more clearly tied together than they are today." That day is dawning now.



© 1982 by SEYBOLD PUBLICATIONS, INC., P.O. Box 644, Media, Pennsylvania 19063. Telephone: (215) 565-2480. Reproduction in whole or in part without express written permission is prohibited. Permission will not be granted to suppliers to reproduce any part of The Seybold Report on Publishing Systems for commercial purposes.

EKC 000142124



The Combiskop, heart of the Chromacom system. This workstation permits page assembly, retouching, color adjustment, and a variety of other operations.

The Hell Chromacom Digital Color System

THE HELL CHROMACOM SYSTEM (marketed by HCM in North America) is a tool for the electronic assembly of full pages of color imagery. The output of the system is screened and color-separated film, ready for platemaking (or, for gravure, since a digital cylinder engraving machine may be driven directly). The key component of the Chromacom system is the Combiskop, a workstation at which scanned-in images can be assembled into pages and a wide variety of adjustments to the individual images and the page as a whole can be made.

For Hell, the Chromacom system is the latest step in an evolutionary series of products. Unlike their most important competitor in this market, Scitex, Hell has a long tradition of providing electronic products to the graphic arts industry. It is a background emphasized both by Hell and by its loyal customers.

Company history

The Chromacom system is manufactured by Dr.-Ing. Rudolf Hell GmbH in Kiel, West Germany. The company was founded in Berlin in 1929 by Rudolf Hell. Dr. Hell was 28, and he had already written a book on the infant technology of television. (Hell and his professor, Max Dieckmann, had made the first public demonstration of wireless transmission of television pictures.) Among the first products offered by the new company were facsimile machines for newspaper use. Facsimile continues to be an important Hell product area to this day. Other early product lines included radio compasses, direction finders, and Morse code recorders.

At the end of World War II, in 1945, Hell ceased operations. It was re-started two years later, in Kiel. The initial activity was repairing facsimile and Morse code recorders, but soon a variety of other tasks occupied the company: restoring the newspaper wire-service network, building a facsimile service for the Post Office, and designing a phototypesetting system. As time went on, Hell concentrated more and more on products for printing and publishing.

Hell's best-known products today—the line of color-separation scanners and related equipment—stem from a 1953 demonstration in which facsimile transmission was used to engrave a printing plate directly, instead of requiring photoengraving as an intermediate step. The initial product, called a "Klischograph," made raised plates for black-and-white letterpress. Color capabilities, and output suitable for offset and gravure, followed later. The direct-engraving technology led to the current "Helio-Klischograph" product, which engraves gravure cylinders directly using diamond styli under computer control. The color scanning technology led to the present "Chromagraph" family of scanners.

Hell also makes equipment for typesetting—the Hell "Digiser" systems, which were introduced in 1965. These have sold well in Europe, particularly in Germany and in Switzerland, Austria, and Yugoslavia. Digisers have been sold in the U.S. market two different times by two different companies. When the Digiser was first introduced, RCA and Siemens had a cooperative agreement (Siemens sold RCA computers under its own name in Germany). The original Hell Digiser was sold in the U.S. as the RCA VideoComp 820.¹ (RCA also sold Hell color scanners.)

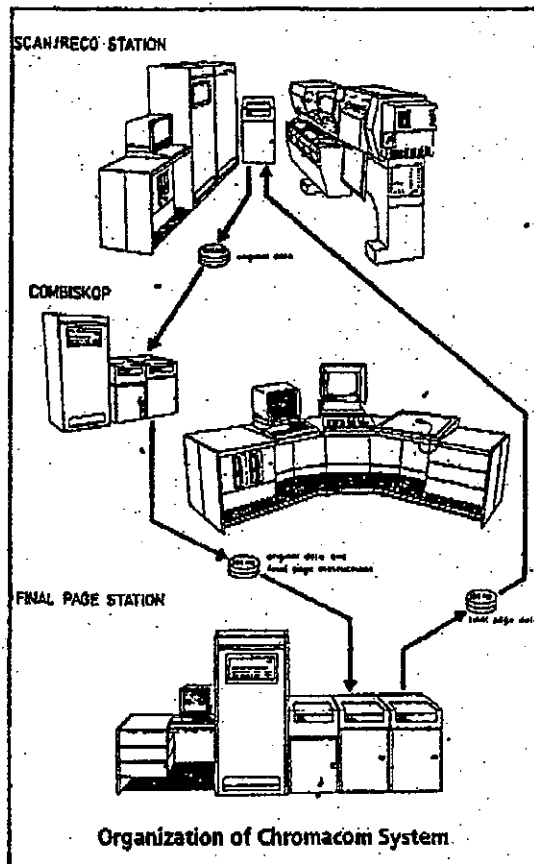
More recently HCM, Hell's North American subsidiary, offered the Digiser 20T typesetter in this market. But the machine was late into the market, higher-priced than its American competitors, and did not offer a full library of U.S. type faces. It has since been withdrawn, but it continues to sell well in Europe.

Products for the printing industry dominate Hell's output. The current annual report does not give figures on prod-

¹Eventually, the VideoComp and the Digiser evolved into completely separate product lines. Later RCA VideoComp models were hybrid machines with an RCA "front-end" (the computer/controller) and Hell "back-end" optical bed. The VideoComp product line was subsequently acquired by Information International. Since then III and Hell have proceeded on their own separate development paths. Current III VideoComps and Hell Digisers share no components in common.

7-4

The Seybold Report on Publishing Systems



uct areas, but a 1979 report showed that over 80% of Hell's sales were printing-related. In 1979, 62% of sales were of scanners, engraving machines, and related equipment; and 19% were phototypesetting-related. Facsimile devices contributed 17% and textile equipment 2%.

Hell is now a wholly-owned subsidiary of Siemens AG. Siemens, the giant European electrical and electronic conglomerate, having held an 80% interest in the company, purchased the remaining 20% from Dr. Hell in 1980. (Dr. Hell, now 81, is honorary chairman of the board.) Hell's sales for 1980/81 were 395 million marks (about \$186 million) with an after-tax profit of DM 22 million. Total sales for Siemens were DM 32 billion in 1979/80.

International scope. Hell is very much an international company, with 72% of total output being exported. About half of Hell's sales are in Europe (including Germany), approximately a quarter are in North America, and about 10% in Japan.

The Chromacom system

Although it continues to sell "straight" scanners, the key to Hell's future clearly lies with the Chromacom digital color system and its related input scanners and output recorders. The basic principles of such a system are by now familiar to most of our readers. Continuous-tone color transparencies or

prints are "read" on a color-separation scanner and recorded onto disk as digitized continuous-tone (unscreened) pictures. The amounts of data thus recorded are immense: each sample point, of which there are typically 300 per inch in each dimension, is represented by 24 bits of data. This means that for each square inch of image area, there are roughly two million bits (or a quarter of a megabyte) of data. This much data cannot be displayed and manipulated in real time with today's technology, so a coarse-resolution sampling is used for operations like color correction, image placement, and retouching, which have to be done interactively. The functions performed by the operator on the coarse data are then repeated on the full resolution data as an off-line process, called "final page processing." The final step is to record the completed page as sets of color-separated negatives. The screening is performed at this time.

The way data is transferred from one process to the next is usually by moving a disk pack from one disk drive to the next, although Hell has recently begun offering a facility which avoids the necessity of doing this.

The Combiskop

The Chromacom system has several "stations," some essential and some optional (the details are provided under "Putting together a system"), but all that is absolutely necessary is a scanning and recording station and a Combiskop station.

The Combiskop is the heart of the Chromacom system. It is the capability of the Combiskop which make it possible to do things with images which could not have been done conventionally. These facilities include page-assembly, color correction, and retouching. The Combiskop is also a key element in the cost-effectiveness of the Chromacom in use. The number of pages per hour that can be run on the Combiskop, and the number of times a given page has to be called back to the Combiskop because of revisions requested after proofing, are likely to determine whether the whole system can pay for itself or not.

Because of its importance, we will describe the operation of the Combiskop in detail.

The operator controls the Combiskop primarily using two input devices: a function box (to indicate which operation is desired) and a digitizing tablet (to indicate positions on the screen). The operator sees the effects of each operation on a video display screen. Movements of the "puck" on the digitizing tablet are reflected in cursor movements on the screen. Most operations involve a single cursor, but for some two cursors are displayed. There is also an alphanumeric VDT at the Combiskop, but this is little used. It is primarily intended for running utility programs. During Combiskop operations, the VDT displays a "job listing" of each function the operator invokes. Error messages are displayed on it.

Image size and resolution. Before describing the functions which are available on the Combiskop, a little bit of background is needed concerning how images are stored and displayed.

The screen of the Combiskop can display a maximum of 512 "pixels" (image points) in the horizontal and vertical directions. An 8" square image, at 300 dots per inch, requires 2400 pixels in each direction to be shown at full resolution. Such resolution is beyond the state of the video-display art,

19. NOV. 2004 10:45

44_01937_5466623

NO. 2507 P. 4

so the Combiskop (or any competing product) must show either the whole image at reduced resolution or a portion of the image at full resolution.

The Combiskop offers both possibilities. It keeps each image on disk in two forms: full resolution and coarse resolution. The latter is computed from the former by taking a center-weighted average of a square block of pixels and using that average to represent the whole block when displaying the image at coarse resolution. For example, the coarse resolution image might have one pixel to represent each group of 49 pixels (i.e., a 7x7 block) in the full-resolution image. In this example, the coarse-resolution image would contain roughly 2% of the number of pixels in the full-resolution version. Depending on the scanning resolution and the degree of enlargement required, the number of full-resolution pixels per coarse-resolution pixel could vary. For instance, line art is scanned at very high resolutions (up to 1800 lines per inch) so the coarse-resolution data for screen display might have only one pixel for each 20x20 block of full-resolution data.

The final size of a piece of art generally needs to be determined at input scanning time, because that is when the two disk versions are created. As will be seen, it is possible—but time-consuming—to change the size of an image at the Combiskop.

Either resolution image can be loaded into image memory and displayed on the screen. If full resolution is chosen, a 512-pixel-by-512-pixel block is loaded. Generally, this is only part of the whole image. Some operations are handled best at full resolution, and these must often be done one piece of the image at a time.

To complicate matters further, there are two other operations that affect the size of an image on the screen. One is the "zoom" feature, which causes apparent enlargement of an image by "pixel replication." This consists of copying each image pixel several times in the vertical and horizontal directions so the size of the image is increased without any new data being introduced. The available zoom factors are 2x, 4x, and 8x. The operator can apply the zoom feature to whatever is on the screen—either full- or coarse-resolution images—and the change is instantaneous. Zoom is provided for operator convenience and has no effect on the underlying image data.

The other operation affecting size is the "rotation and scale change" function. This is a function which can be used to change both the size and orientation of an image. It is relatively slow. Depending on the size of the object to be rotated, it may occupy the system for several minutes after the operator has indicated the desired size and orientation. This is the only operation which can actually change the size at which an item will be output from the size specified during scanning.


We note that Scitex now has interactive on-screen sizing and rotation facilities on its system. This is, we think, a very useful feature and one which Hell should add. To do so, Hell would need to use special-purpose hardware, designed for the task, just as Scitex does.

Image memories. The Combiskop can seem complicated at first glance. But the principles of working with it are simple. The key technical notion that is required in order to understand the Combiskop and its procedures is the idea of an "image memory." This is an area of computer memory set aside for storing pictures temporarily while they are being worked on.

For the most part, images are stored on disk. But when they are to be displayed, they are called into the image memories of the Combiskop. There are two image memories, so two different images can reside in the Combiskop at one time. The operator can select the image in either memory for display on the screen. The normal page-assembly procedure is to call each image in turn from disk into memory two; perform operations like retouching, color adjustment, and mask creation; then add it to the page which is being assembled in memory one. After the page is completed, it is written back to disk from memory one.

Masks. Although they cannot be seen in the finished page, masks are a fundamental part of color image assembly, both in conventional processes and on electronic systems like the Chromacom. Masks isolate an image area for subsequent manipulations. They can be thought of as "windows" of arbitrary shape through which a specific image can show.

For example, if the operator wants to pick-up an item (a tape recorder, for instance) out of a scanned-in image and put it into a page which is being assembled, he can use a mask



The page-creation process. This simplified demonstration page begins with a mask (left) and a page-size color vignette (left center). A border is added, one piece of art is placed within the border, and a window for a second piece of art is created (right center). The second piece of art is positioned in the window, completing the page (right).

19. NOV. 2004 10:46

44_01937_5466623

NO. 2507 P. 5

7-6

The Seybold Report on Publishing Systems

which blocks out the unwanted background from the tape-recorder picture, and leaves only the recorder itself showing through. When the mask and the picture are placed together on the page the mask permits the existing page to display only to the extent that it is outside of the area occupied by the recorder. Anything which was in the area now occupied by the recorder is hidden. This is an example of what Hell calls "foreground" masking, since the mask causes the new image to be placed "on top of" the existing page.

Another type of masking, "background" masking, occurs when a mask is positioned on a page and an image is positioned behind it. For example, suppose the layout calls for the tape recorder to appear in a framed box inset into the upper left-hand corner of the page. In this case, a rectangular mask and a computer-generated frame would be created first in the proper position on the page. Then the tape recorder image would be called to the screen. The mask would function as a window through which part of the tape recorder image could be seen. Moving the image freely with the digitizer "puck," the operator would position the tape recorder within its stationary frame. No part of the page outside the frame would be affected. The new image appears to be "in back of" all the existing page elements.

A final type of masking, which is called "mixed ground," involves two masks. One mask, which is stationary, protects parts of the existing page from being covered by the new image. The other mask, which moves with the new image, causes it to cover unprotected areas of the page. Thus, as the new image is moved around the partially-assembled page, it will disappear behind some objects and hide others from view.

Creating masks. The importance of masks should be clear by now. Quite a bit of what the Combiskop operator does is related to creating and using masks.


There are three ways to create masks at the Combiskop. The easiest type of mask to create is the machine-generated geometric shape. The Combiskop will automatically generate rectangles, circles, ellipses, and polygons from points input on the digitizing tablet.

If the background of an image is to be masked out, and if that background is of fairly uniform color, there is an automatic masking function which can create a mask covering all areas which are of that color (or very close to it). If the background is non-uniform, or if the image of interest contains areas of the background color, this method is less useful.

The final method of creating a mask is simply to draw it on the screen, using the digitizing tablet. This process, called "contouring" by Hell, must be done at full resolution. Buttons on the digitizer's puck permit either creation or erasing of mask areas.

After a mask has been created, it can be stored in any of seven mask memories of the Combiskop. It can also be written to disk for future use. If a mask which has been stored on disk is needed again but in a slightly different form, it can be called back to the screen for editing.

Using the Combiskop. The operator tells the Chromacom what tasks to perform via a function box which has a key for each function. The keys of the function box are in five groups. In the middle is a numeric keypad. At the left are a group which controls the cursor and digitizing tablet and a



The Combiskop function keyboard. The keys are in functional groups. The four knobs are for color adjustments. Below the word "Combiskop" is a single-line display showing the last keys pressed.

group which controls the display and placement of images and masks. At the right are a group which cause masks and frames to be generated, and a group involving generation or correction of color.

The layout of the function box is sensible and seems reasonably easy to learn in spite of the large number of keys (there are about 100 keys). It may seem like a small point, but we are surprised that Hell does not provide English-language abbreviations on the keys for systems sent to English-speaking countries. Some of the German abbreviations are close enough to the English equivalents to be useful (e.g., "KEL" for "create ellipse"), but many are not ("FWD" for "define color value," "BDM" for "rotate and scale image"). We think it would be very desirable to change these abbreviations, as well as those in the job listing (see below) to reflect their meanings in English.


The commands in the first function-box key grouping (cursor and digitizer commands) are basically set-up commands. They don't cause any modification of images. This group of commands includes selection of the shape and color of the cursor, the choice of which cursor to move next (when both are being displayed), and adjusting the positions of the on-screen cursor and the "puck" of the digitizing tablet so that the layout corresponds properly with the image on the screen.

The next group is the image and mask-manipulation commands. These include such things as loading images from disk (either at full or coarse resolution), loading masks from disk, positioning images and masks on the page, and changing the zoom factor. Also in this group is the auto-mask command, whose use was described earlier.

The next group of keys controls the creation of machine-generated frames and masks. Rectangular, circular, elliptical, and polygonal masks and frames can be created. Frames can be of any specified thickness and can have rounded, angled or square corners. A frame may be created so that it lies along the inside of the border of a mask, along the outside of the border, or so that it straddles the border. Also in this group are the contouring commands for drawing masks freehand. Two additional functions that are in this group are those which cause rotation and enlargement of images and masks.

The final group of functions are those concerned with color manipulation. One set of vignette-defining keys permits the operator to enter specific color values at selected points and the system generates a vignette that incorporates those colors at those points. Another set is used for retouching. The operator can select the size, shape, color, and speed of action of the electronic retouching "brush." Another set applies conventional color corrections to the highlights, middle tones, or shadows of an image. All of these functions may be

EKC 000142128



Working with scanned-in type and rules. Left: This magazine cover demonstrates working with line art. Blocks of type were given various solid colors and vignettes. Top right: The job displayed on the Combiskop screen. Bottom right: The colored type is "spread" under the dark background, but where it emerges onto the white paper it is not. (Yellow separation negative, 200%).

applied to the entire image, or their effect may be restricted to an area defined by a mask.

The job listing. As each function key is pressed, its abbreviation appears on the alphanumeric VDT associated with the Combiskop. There is usually not much need to refer to this "job listing" during normal operation, but it does serve an important function if the job needs to be rerun in a slightly different form at a later time, or if the operator discovers, part way through a job, that some step had been omitted earlier in the page-assembly process.

In situations like these, the job listing can be rerun, providing a kind of "instant replay" of the operations that have been done on the Combiskop. The only things which have to be re-done are those (such as retouching) which involve cursor movements during the operation. Furthermore, the job listing can be edited. This means that if two jobs differ in only a few details, the job listing from the first may be edited to produce a listing which will cause the second to be run

automatically. We'll return to the job listing in describing the Layout Programmer station.

Final page processing

After all operations at the Combiskop have been completed, and before a page can be output, there is a step called "final page processing" through which the image data must pass. When a page is assembled on the Combiskop screen, most operations are carried out at coarse resolution. The full-resolution data still resides on disk in its original form. Two things have to happen before the output process can begin: all of the retouching, rotating, masking, and other image-altering steps which were done on the Combiskop have to be applied to the full-resolution data; and the various pieces of the image have to be sorted into the raster sequence in which they will be output.

This process can be lengthy. In some cases, it may even exceed the time it took the Combiskop operator to assemble

7-8

The Seybold Report on Publishing Systems

the page. This process can be done on the Combiskop, but since it involves only the minicomputer and not the operator's console, and since the Combiskop cannot be used for normal operations when it is running the final page process, it generally doesn't make sense to use the Combiskop for this process. For this reason, most Chromacom purchasers buy a "final page station"—an extra minicomputer with a pair of its own disk drives dedicated to performing just this operation.

Putting together a system

Chromacom jobs go through three major steps: input, assembly, and output. Input and output are handled by what Hell calls the "Scan/Reco" (for scanning and recording) station, and assembly is done at the Combiskop. The minimum configuration consists of just these two stations.

Scan/Reco. The Scan/Reco station can use a standard Hell scanner (the DC 350 or the large-format CP 340) for both input and output. Or one of the new output recorders (see below) can be used, leaving the scanner to perform input only. The DC 300B can be used as an input scanner but not as a recorder. In any case, there will be a Siemens minicomputer with an operator's VDT and at least two 300-MB disk drives. A scanner being used as a Chromacom input and output device can still be used normally as a stand-alone scanner when it is not needed for use with the system.

Combiskop. The Combiskop station consists of the Combiskop itself (including the display electronics, the digitizing tablet, and two floppy-disk drives) and a second Siemens minicomputer with VDT and 300-MB drives. At this station, there would also be an 80 MB drive for software and stored job files.

Final page station. Very few customers would choose this minimum configuration, however. There would almost always be a third minicomputer-plus-disk station at which final page processing would occur. (The alternative would be to run the final page processing on the Combiskop station, but this would make it unavailable for normal operations for periods of 15-30 minutes or more per page.)

Output recorders. Hell offers several alternative output devices for the system. The standard Hell scanners have already been mentioned. In addition, there are two output recorders and a proof-recorder. The CR 401 automatic recorder can handle film up to 21" x 29". It loads its own film, exposes it automatically, and deposits it in an output cassette or straight into a film processor. It doesn't require a darkroom. The CR 402 is a large-format recorder that can handle 44" x 50" film (the same as the CP 340 scanner). It is hand-loaded and must be operated in a darkroom.

The CR 403 proof-recorder is an interesting new product, introduced at DruPA. It handles 21" x 29" color film or paper on which an unscreened proof can be recorded. It produces its images using two lasers. One is a Helium-Neon laser with output in the red portion of the spectrum. The other is an Argon laser with both blue and green components in its output. The blue and green portions are optically separated, making three beams in all. The three beams

are independently modulated to expose the full-color image on photographic film or paper.

Layout Programmer. Another option which could speed workflow in a heavily-loaded system is the "Layout Programmer." This workstation offers the same frame and mask generating options as the Combiskop, but it doesn't handle scanned images and the various image-related functions (color-correction, airbrushing, etc.) are not available. The bulk of the work of many jobs, however, can be handled with just the functions that are available on the Layout Programmer. The Layout Programmer does not show the true colors of frames and tint areas. It can display only eight different colors. However, the operator can specify the color value that each item will have when it reaches the Combiskop.

The operations at the Layout Programmer are recorded on floppy disk as a job listing. This is subsequently "played back" on the Combiskop. At each point where the Combiskop operator needs to intervene, the job listing will include a "pause" command. The Combiskop operator can then make whatever corrections or adjustments the job calls for and resume running the job. (Further enhancements to the Layout Programmer are in the works, as described in "Plans for the Future.")

Data transfer. A high-speed magnetic tape facility (6250 bits per inch, 75 inches per second) is available for archiving and for transferring image data to other sites.

Each Chromacom "station" is essentially a stand-alone subsystem with its own minicomputer and disk drives. Within the Chromacom system, the normal way to move image data from one "station" to the next is to stop the disk drive on one station, remove the disk pack, and install it on the drive for the next station. This process is not good for disk packs or drives (both of which fare better if they are turned on and left) and it is an operational nuisance to be constantly moving disk packs around.

Hell has announced a data-switching facility which addresses this important problem. With the data-switching "network" option, a given disk drive can be connected to any of the stations without removing the disk packs. Thus, as a job moves from input scanning to Combiskop to final-page processing to output, it remains on the same disk drive but that drive is connected to each station in turn. The switching facility comes in three "levels." The minimum level provides a manually-operated switchbox which performs the switching function and nothing more. The second level provides automatic switching under the control of a minicomputer. The third level provides additional software on the switching minicomputer for such things as job tracking, cost estimation, and system-wide file management. These facilities will give Hell a file-management capability similar to Scitex's.

North America: HCM Graphic Systems, Inc.
300 Rabro Drive East
Hauppauge, NY 11788
Telephone: (516) 582-6520

World headquarters: Dr.-Ing. Rudolf Hell GmbH
Grenzstrasse
2300 Kiel 14, West Germany
Telephone: (04 31) 2 00 11

Flat-bed scanner. In early 1984, Hell will begin deliveries of a high-speed black-and-white flat-bed input scanner, the CN 420. This scanner can accommodate a broadsheet newspaper page (the scanning area is 19" x 23½") and is derived from Hell's products for the newspaper facsimile market. It will be able to handle both transparent and opaque originals. Resolutions up to 2500 lines per inch will be supported. At 1830 lines per inch, it can scan at a rate of 4.6 seconds per inch, which is equivalent to a broadsheet page in just over a minute and a quarter. This scanner will be useful for scanning type, line art, and pre-screened halftones.

Architecture of the Combiskop

The apparent simplicity of operation of the Combiskop belies its complex architecture. The interactivity of the display is the result of some very sophisticated image-processing that goes on continuously inside the Combiskop.

There is a Siemens minicomputer associated with the Combiskop, as there is with each station of a Chromacom installation. But the minicomputer provides very little of the processing power that is resident in the Combiskop. The main image-processing capability resides in a special-purpose display processor which Hell buys from the DeAnza division of Gould Corporation. This processor is under the control of a DEC LSI-11/23, which, in turn, is connected to the Siemens minicomputer. The 11/23 accepts data from the function box and digitizer and handles the floppy disks, as well as giving the DeAnza unit its instructions.

The DeAnza display processor's full-time function is to keep the color display running. In the process of doing this, it recomputes the color value of every one of the quarter-million pixels on the screen every thirtieth of a second. The processor gives each pixel a 24-bit value. Eight bits each are used to define the red, green, and blue components at each point. There is enough memory for two such 512-by-512-by-24-bit-deep images, plus a third temporary area of the same size used during retouching and other image alterations.

There are also eight "overlay" areas, each 512 by 512 by one bit. These are the basis for the Chromacom masks.

The image processor is constructed so that the screen is constantly refreshed by reading the entire contents of one of the image memories every thirtieth of a second. A number of processes, such as color adjustment and image shifts in the vertical and horizontal direction, can be done "on the fly" in the circuits between the image memory and the video tube. This gives the Combiskop its fluid interactivity for these processes. The "zoom" feature is also handled by this hardware, as are the two cursors.

The DeAnza processor can also produce displays from data which is partly being read from one image memory and partly from another. One of the overlay memories acts as a "switchbox" for the processor. In any position where the overlay memory contains a "one" bit, the data from one image memory is used. Positions where the overlay has a "zero" bit are read from the other memory. This feature underlies the masking capabilities of the Combiskop.

The raw processing power of the display processor is awesome. It is best appreciated by considering the difference in response times between functions like moving a picture around on the screen or changing the zoom factor (which are both instantaneous) and generating a vignette (a number of

seconds) or rotating a sizable image (a number of minutes). The latter functions are done by software in the LSI-11/23, the former by the DeAnza hardware.

In fact, we think that Hell needs to push DeAnza for hardware to support real-time sizing and image rotation. This is one area where the current Scitex system offers a clear advantage over the Chromacom. Scitex's initial offering suffered from the same problem, but a hardware solution has since been found.

The Chromacom in the field

To get an understanding of what the Chromacom system means to its users, we visited two user sites. They were remarkably different in their approach to the system. One user, Kwik International in New York City, had been using the system for over two years. The emphasis at Kwik was on fast-turnaround advertising work. Kwik's history of color-separation work goes back many years before the Chromacom. Kwik has impressive facilities for all types of prepress work, both conventional and electronic. They were shown to us by Kwik's president, Dan Sirota.

The other user was Time-Life Books in Alexandria, Virginia. The system there was brand new and not yet in use for real production jobs. The system was purchased for in-house use, primarily in the production of high-quality "coffee-table" books. Time-Life Books had no prior experience with in-house color separation before purchasing the Chromacom. Tom Boynton, project manager, showed us around the carefully designed and appointed facilities.

A question we put to both Sirota and Boynton was their reason for selecting the Hell system over the competitive Scitex offering. Both men had obviously been asked the question many times before. They gave several reasons, but the most important one seemed to be the ability to get the entire package, including maintenance and support, from a single vendor. Scitex does not make an input scanner, so Scitex installations inevitably involve multi-vendor support.

Both Boynton and Sirota felt that type and line-work should normally be handled separately on film and not be run through the system. Boynton pointed out that Time-Life



Dot-etcher checking negatives. Dot-etching is one of several labor-intensive steps which the Chromacom system bypasses. This photo was taken at Kwik International.

19. NOV. 2004 10:49

44_01937_5466623

NO. 2507 P. 9

7-10

The Seybold Report on Publishing Systems

books are often published in several languages, but with the same layout and pictures. If the type is on separate negatives, then only those negatives need to be re-done when a book comes out in translation. If the type were handled on the system, it would appear on the same negative with the black separation.

Sirota gave a different reason for keeping the type separate. Eighty percent of changes called for at the proof stage are changes to the text. These changes are easily handled conventionally if the type is on a separate piece of film. It would be inefficient to go back to the Combiskop to make simple wording or price changes in an ad. Sirota noted, however, that there are times when it makes very good sense to scan type and line art. This is often the case, for example, when tints or vignettes involving line work are called for. Kwik

had recently taken a job which required that a black-and-white engineering drawing with lots of fine lines had to be reproduced as a color job with the lines being one tint and the background another. It would have been extremely difficult to do using conventional masking and stripping techniques, but it was easy on the Combiskop.

On the subject of cost-justification, Boynton had some very specific projections concerning the Time-Life installation. He said that the current annual cost of outside color-

(Text resumes on page 15.)

The insert which follows this page was provided by HCM to illustrate some of the capabilities of the Chromacom system.

Chromacom user list

As of this writing, the Chromacom has been installed in North America at the following thirteen sites:

Bonac-Batten
Toronto, Ontario, Canada
Gravure Systems
Florence, Kentucky
HCM Demonstration Studio
Los Angeles
HCM Demonstration Studio
New York, New York
Kwik International
New York, New York
Lahigh Electronic Color
Chicago, Illinois
Lanman Lithography
Orlando, Florida
MacLean Hunter Publications
Toronto, Ontario, Canada
Pacific Lithograph
San Francisco, California
R.L. Donnelley
Chicago, Illinois
Spectrum Incorporated
Minneapolis, Minnesota
Time-Life Books
Alexandria, Virginia
Weston Engraving
Minneapolis, Minnesota
Seven more systems have been sold, but not yet installed, in the following locations: Boston; Los Angeles (two systems); Long Island, New York; Philadelphia; Portland, Oregon; Toronto, Ontario.

In Europe, there are 53 Chromacom installations as of this writing.

Adplates Ltd.
London, England
Angsa Lito AB
Stockholm, Sweden
AS Clide
Oslo, Norway
Burda GmbH
Offenburg, W. Germany
Cino De Duca
Blois, France
Cino De Duca
Maison Afort, France
Cosp Offset
Montreuil, France
D. S. Colour International Ltd.
London, England
De Back & Paulich
Wetteren, Belgium
Graafmestudio
Helsinki, Finland
Francis Imprimerie
Ozoir-la-Ferrière, France

Hell-Studio
Kiel, W. Germany
Hellelectron f. Studio
Stockholm, Sweden
Helsingvln Kuvataitutehdas Oy
Helsinki, Finland
HTF Scanner Team
Krefeld, W. Germany
Ite
Turin, Italy
Interrepro
Münchenstein, Switzerland
Köhn Repro
Copenhagen, Denmark
Krammer
Linz, Austria
Kunnallispaino
Vantaa, Finland
L. Europe
Brussels, Belgium
Leleux
Brussels, Belgium
Laudert & Co.
Vreden, W. Germany
Malmö Repro-Kopia AB
Malmö, Sweden
Mayday Reproductions Ltd.
London, England
Mohndruck
Götersloh, W. Germany
Mondadori
Verona, Italy
Nefflex
Zug, Switzerland
NEFLI
Haarlem, Netherlands
Neue Chemiegraphie AG
Zürich, Switzerland
Nureg GmbH
Nuremberg, W. Germany
Oestricher & Wagner
Munich, W. Germany
Otava
Helsinki, Finland
Persike f. Studio
Mitham, England
Pesavento & Co.
Zürich, Switzerland
Photomatic
Lyon, France
Prolith AG
Köniz, Switzerland
Promograph S.A.
Madrid, Spain
Repro Zentrum
Klagenfurt, Austria
Reprostudy S.A.
Hosp. de Llobregat, Spain
San Paolo
Alba, Italy

Seibald Druck & Verlag
Nuremberg, W. Germany
Schaurfeiberger AG
Winterthur, Switzerland
Schaffler
Frankfurt, W. Germany
Schmidt
Stuttgart, W. Germany
Schmidt-Repro
Dornbirn, Austria
Siemens f. Studio
Milano, Italy
Sierstiens f. Studio
Paris, France
Siemens f. Studio
Stuttgart, W. Germany
Süddeutsche Klischee-Druck
Munich, W. Germany
Tessa
Brussels, Belgium
TGI
Glanerbrug, Netherlands
Time Scan
Leinfelden, W. Germany
Vau Velle Photo Lito
Leeds, England
Wirth f. Studio
Frankfurt, W. Germany
WWS Repro
Ditzingen, W. Germany
Zeno GmbH
Münster
Zilling f. Studio
Neuss, W. Germany
Zuliani S. A.
Montreux, Switzerland

In Asia, Australia, and Africa there are these installations:

Curren
Sydney, Australia
Hirt & Carter
Capetown, South Africa
Kaigai f. Studio
Tokyo, Japan
Koei Insatsu
Japan
Mika Seihan
Tokyo, Japan
Phetra f. Studio
Johannesburg, South Africa
Scanagraphix
Melbourne, Australia
Sennelisha
Tokyo, Japan
Show Ads
Melbourne, Australia
Siemens f. Studio
Melbourne, Australia

EKC 000142132

19. NOV. 2004 10:49

44_01937_5466623

NO. 2507 P. 10

separation services is about \$3 million and that the on-going costs of the Chromacom system will be \$1.2-\$1.5 million per year. Boynton mentioned as important in keeping costs down the need for a single decision-maker for approval purposes. If several people are involved in the approval process, it may take many proof cycles to please them all. That would eat up most of the savings.

Sirota emphasized that Kwik had had to go through a learning process before it learned how best to make the system pay for itself. He pointed out that the learning process involved not just the system operators, but also the sales force and the company management. To make money with the Chromacom, Sirota says, "management has to know as much as the operator." Cost estimating for the system is a key area where this is true. While Kwik's sales force can make estimates for conventional jobs, all estimating for the Chromacom is still done centrally. Training the sales force on all the ins and outs of the Chromacom would simply be too hard, and there are no easy rules of thumb.

Sirota's shop is set up to make just about any type of proof a customer might want including press proofs. Boynton is standardizing on Chromalin proofs. Both men emphasized the variable nature of the printing process as limiting the value of proofs. Sirota said he had once sent out the same set of separations to five different printers, all of them highly regarded and all equipped with the proper densitometric equipment to presumably match specifications exactly. The results were surprisingly non-uniform. Boynton suggested that proofing technology was forcing printers to do a better job. He said that prior to the widespread use of proofs for on-press quality control, printers weren't required to meet objective standards. The implication was that Chromalins and similar proofs, properly made, constitute a more reliable basis for judging an image than press proofs, which may actually be less dependable in representing the final printed piece.

Operator qualifications. One of the areas of divergent opinion between Boynton and Sirota was the suitable background for potential Combiskop operators. Boynton downplayed the need for prior graphic arts experience. He said that a sense of humor was the key requirement—learning new skills and breaking in new equipment is always a trying experience and a sense of humor is important in dealing with it. Boynton's second consideration was intelligence, with graphic arts experience ranking a distant third. Boynton means what he says: one of his Combiskop operators had been a typist prior to the purchase of the system.

Boynton does not see the Combiskop as a place where color adjustments should be made (a fact that helps to explain why he feels graphic arts experience is not too important). If the scanner is set up correctly, he expects that an image can be passed through the page-assembly process and output without correction. If the proof shows a need for color adjustment, then the image can be brought back up on the Combiskop and adjusted as specified from the proof. Boynton feels that primary responsibility for color control should rest with the input scanner operator (who, Boynton feels, must have color separation experience) and with the quality control department.

Dan Sirota also stresses intelligence as a key trait for Combiskop operators, but he believes that a thorough

grounding in graphic arts practice and the underlying theory are very important. As a result, the operators at Kwik have solid backgrounds in conventional color work.

Both Sirota and Boynton see the technology of the Chromacom system opening up new areas of endeavor. Boynton stresses cost savings. Shorter print runs will be possible because pre-press investment can be recovered on a smaller sales volume. He sees that Time-Life will be able to offer books tailored to smaller audiences. As a facetious example, he suggests the title "Plumbing for the Left-Handed."

Use for designers. Sirota emphasizes the new creative possibilities. There are many things that the Chromacom can do which would be impossible or prohibitively expensive using conventional techniques. He sees signs that ad agency art directors are beginning to plan jobs with such possibilities in mind. Some art directors who want to use the capabilities today are holding back, Sirota says. They want to have other shops to fall back on in case Kwik's system is down or overbooked when they need a job produced. Sirota is not too chagrined at the prospect of other systems being installed in New York, since those installations mean that more art directors will feel easy about planning jobs that involve the system.

Sirota foresees a day when designers will be able to work on a system of this type. He relates the story of a clothing designer who came into the Kwik facility to check on an ad. Intrigued by the Combiskop, he asked for several changes to be made to the suit that was pictured on the display: the color was changed, the lapels were made narrower, the vest eliminated, the shoulders rounded, etc. Finally, satisfied, he announced that he was going to go back to his shop and create the suit he had just seen on the monitor.

Economics: cost-justifying the system

There seem to be two major applications for the Chromacom system. It can be viewed primarily as a way of automating conventional stripping, or it can be viewed as a device for special effects which are difficult to obtain by other methods.

If it is viewed primarily as a stripping tool, then the way to cost-justify the system is to push as many pages as possible through the system. In this case, it is important not to spend too much time on retouching, color spotting, and other niceties. With all the facilities that the system puts at the operator's disposal, it is tempting to fix up problems that would not be corrected in conventional processing. But unless such work has been allowed for in pricing the job, time spent on such activities is non-revenue-producing time.

This approach to using the system has several advantages. Jobs can be estimated and thought of by the customer and the sales force as if they were to be handled conventionally. There would be relatively little retraining involved in those areas. The sales effort could emphasize jobs with lots of pages, to keep the Chromacom system busy. Many successful Chromacom users have taken this approach.

Here is how the cost-justification might be achieved, using figures provided by HCM. Suppose a shop with a Chromacom could produce 300 pages per month at a selling price of \$500 per page. This could be done with two shifts, according to HCM, providing the work is not primarily ads.

Revenues would be \$150,000 per month. Against this figure must be balanced the costs. Labor would be roughly

19. NOV. 2004 10:50

44_01937_5466623

NO. 2507 P. 11

7-16

The Seybold Report on Publishing Systems

\$20,000 per month. The interest on a seven-year loan which paid for the Chromacom would also be about \$20,000, and payments of principle would be another \$24,000. There would be roughly \$6,000 per month for the service contract. Beyond these items, which total about \$70,000 per month, there would be materials, sales costs and various overheads, but it is evident that the payback could be quite attractive if these figures are realistic.

An alternative approach to using the Chromacom, and one that makes use of the real power of the system, is to concentrate on work that is difficult by conventional processes but easy on the Chromacom. In this approach, the sales force has to be taught the special advantages of the system for various types of work, and they have to seek out those jobs which are most appropriate. These will often be ads, with relatively few pages but a high price per page. Job estimating is of critical importance, since if the time required on the system is badly underestimated the job is unprofitable, and if it is badly overestimated the job may be lost to a shop using conventional processes.

The special needs of gravure

Gravure printing is a very specialized field. It is noted for its ability to produce high-quality color work without the consistency problems of offset. It is a very attractive approach to printing, except for the difficulty of the pre-press phase.

Gravure costs are dominated by the cost of preparing the huge, ungainly printing cylinders. The cost is so great that only very long press runs can be considered. Runs in the millions of impressions are common, and jobs must generally be at least in the hundreds of thousands to be economically produced by gravure. General-circulation magazines, direct-mail pieces, and Sunday newspaper magazine sections are examples of work which is often done by gravure.

Hell has been in the forefront of automating the production of gravure cylinders. The Hell Helio-Klischograph is a computer-controlled multi-headed engraving machine which engraves cells into copper-coated gravure cylinders with diamond styli. It has been widely accepted by gravure printers. Its only significant competition is the laser-engraved plastic-coated gravure cylinder developed by Crosfield. The first installation of that system is at Sun Printers in England.

Hell recently made public research work on a new engraving process which may form the basis for Hell's gravure products five or six years from now. The new method involves engraving a conventional copper cylinder with an electron beam. This exotic process must be performed in a total vacuum. The process promises two key advantages over present engraving methods: it will be an order of magnitude faster, and it will produce better-quality type and line art.

The latter advantage is due to the fact that with the electron beam, cells need not be placed precisely in a straight line (as they are with the Helio-Klischograph). The electron beam is readily deflected a small amount to either side to accommodate the needs of line art, whereas with the Helio-Klischograph, line art has to be fit to the machine's rigid raster causing the type and line art to have a slightly ragged look.

The benefits of this new technology, if it can be brought to market, will make gravure much more competitive than it is today in terms of smaller print runs and high-quality line art to match the quality of the process color.

Plans for the Future

The following statements, provided to us by HCM for this article, describe the approach Hell/HCM intends to take in developing two new capabilities: merging typeset text with graphics, and providing a pre-Combiskop page-composition station.

Text and graphics. The DC 350S and CP 340S are currently able to scan text (or any line art) in a special high-resolution mode (six times normal). The type or line art thus scanned is then merged with the Combiskop-created geometric figures and both are processed internally at this high resolution. For users with high-volume type requirements, we will soon have a special Raster Image Processor (RIP) that will output type face image data in the standard Chromacom raster format. We will be able to interface this RIP to any front-end system, and we are making arrangements with major American vendors of digital font libraries to license their fonts to our customers. We are planning to have this product available by the middle of next year.

Pre-Combiskop station (Designer station). We are developing an extended version of the Layout Programmer Station that will be able to handle full-color images as well as geometric figures and frames. This station will work with a library of video-format pictures that have been input through a standard television camera, and the originals of the selected pictures will then be scanned in at the regular scanner. The initial hard-copy output at this station will be a monochrome representation of the composed page, which can then be used as proof-copy for approval by an art director or as a layout by the Combiskop operator. Ultimately, this station will have multiplexing and networking capabilities with the Combiskop, the Scan/Reco station, the type RIP, and with others of its own kind. We plan to release more information on this product next year.

Conclusions

The Chromacom has, by now, proved itself a worthy contender in the color page-assembly arena. Chromacom sales are going very well at the moment (at the same time that Scitex has experienced several quarters of flat sales) and the future looks promising.

With the exception of real-time image rotation and sizing, all the basic tools are in place for efficient production, and the plans for future offerings sound appealing. We think the decision to interface to various front ends and to use fonts from various sources (instead of relying entirely on the Hell Digiset fonts, which Hell must have been sorely tempted to do) is a very good one. This will make the Chromacom attractive to a new and extensive market: potential customers with an existing investment in editing and typesetting equipment and a need for color page-assembly.

We are also attracted to the idea of the video-resolution pre-Combiskop station that Hell plans to offer. In some respects, this product sounds similar to the Scitex "Vista" console, introduced at DRUGA. But we like the Hell approach of working with full-color imagery from a television camera. This type of product could be the forerunner of workstations

EKC 000142134

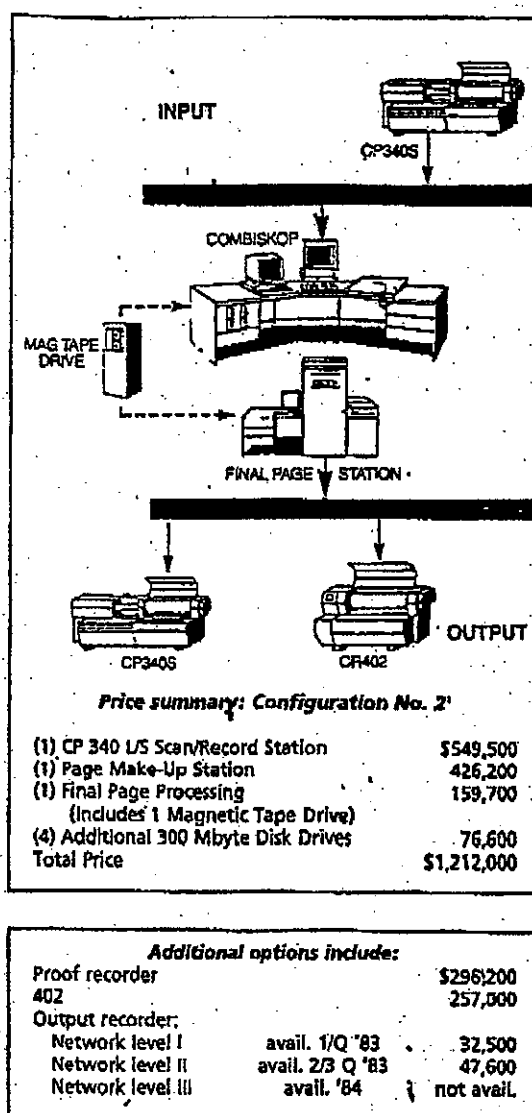

19. NOV. 2004 10:51

44_01937_5466623

NO. 2507 P. 12

The Seybold Report on Publishing Systems

7-17



which are design tools, rather than production ones. Just as the newroom terminal changed the typesetting world, bringing control into the hands of the author/editor and making production more efficient, so a design workstation could change the world of color pre-press, bringing the same kinds of control and efficiency.

In the nature of these new offerings (as well as in the willingness to announce them while they are still under development) we see signs of increasing responsiveness of Hell to the North American market. This we applaud. Hell has, at times in the past, appeared to us unresponsive to (or unconcerned with) the particular requirements of potential customers on this side of the Atlantic, but this appears to be changing.

Hell's unique position as a vendor of both typesetting and color page-assembly systems means that as these areas merge, Hell is very well positioned to maintain a leadership

role. Of the various technologies involved in getting straight to the printing plate from raw inputs, the only one Hell has no announced product for is the ability to make lithographic plates directly from Chromacorn output. But Hell has shown products pointing in this direction in the context of its newspaper facsimile work, and we would expect this last capability to be added in due course.

Hell has a lot of strengths. If the company continues to listen to the needs of its customers, especially when it comes to the American market where much of the action will certainly be in the near future, it should continue to prosper. HCM is making important contributions in product definition and refinement and we expect HCM will be able to play an increasingly important role in Hell's future offerings.

George A. Alexander

EKC 000142135

Second Declaration of George T. Ligler

Exhibit 10

CONFIDENTIAL EXHIBIT

Second Declaration of George T. Ligler

Exhibit 11

CONFIDENTIAL EXHIBIT

Second Declaration of George T. Ligler

Exhibit 12

CONFIDENTIAL EXHIBIT

Second Declaration of George T. Ligler

Exhibit 13

Part1 - Part2 - Single Page

Top Document: JPEG image compression FAQ, part 1/2
Previous Document: [3] When should I use JPEG, and when should I stick with GIF?
Next Document: [5] What are good "quality" settings for JPEG?

[[Usenet FAQs](#) | [Search](#) | [Web FAQs](#) | [Documents](#) | [RFC Index](#)]

[4] How well does JPEG compress images?

Very well indeed, when working with its intended type of image (photographs and suchlike). For full-color images, the uncompressed data is normally 24 bits/pixel. The best known lossless compression methods can compress such data about 2:1 on average. JPEG can typically achieve 10:1 to 20:1 compression without visible loss, bringing the effective storage requirement down to 1 to 2 bits/pixel. 30:1 to 50:1 compression is possible with small to moderate defects, while for very-low-quality purposes such as previews or archive indexes, 100:1 compression is quite feasible. An image compressed 100:1 with JPEG takes up the same space as a full-color one-tenth-scale thumbnail image, yet it retains much more detail than such a thumbnail.

For comparison, a GIF version of the same image would start out by sacrificing most of the color information to reduce the image to 256 colors (8 bits/pixel). This provides 3:1 compression. GIF has additional "LZW" compression built in, but LZW doesn't work very well on typical photographic data; at most you may get 5:1 compression overall, and it's not at all uncommon for LZW to be a net loss (i.e., less than 3:1 overall compression). LZW *does* work well on simpler images such as line drawings, which is why GIF handles that sort of image so well. When a JPEG file is made from full-color photographic data, using a quality setting just high enough to prevent visible loss, the JPEG will typically be a factor of four or five smaller than a GIF file made from the same data.

Gray-scale images do not compress by such large factors. Because the human eye is much more sensitive to brightness variations than to hue variations, JPEG can compress hue data more heavily than brightness (gray-scale) data. A gray-scale JPEG file is generally only about 10%-25% smaller than a full-color JPEG file of similar visual quality. But the uncompressed gray-scale data is only 8 bits/pixel, or one-third the size of the color data, so the calculated compression ratio is much lower. The threshold of visible loss is often around 5:1 compression for gray-scale images.

The exact threshold at which errors become visible depends on your viewing conditions. The smaller an individual pixel, the harder it is to see an error; so errors are more visible on a computer screen (at 70 or so dots/inch) than on a high-quality color printout (300 or more dots/inch). Thus a higher-resolution image can tolerate more compression ... which is fortunate considering it's much bigger to start with. The compression ratios quoted above are typical for screen viewing. Also note that the threshold of visible error varies considerably across images.

AX202069

Top Document: JPEG image compression FAQ, part 1/2

Previous Document: [3] When should I use JPEG, and when should I stick with GIF?

Next Document: [5] What are good "quality" settings for JPEG?

Part1 - Part2 - Single Page

[[Usenet FAQs](#) | [Search](#) | [Web FAQs](#) | [Documents](#) | [RFC Index](#)]

Send corrections/additions to the FAQ Maintainer:

jpeg-info@uunet.uu.net

Last Update June 15 2004 @ 00:27 AM

AX202070

Second Declaration of George T. Ligler

Exhibit 14

CONFIDENTIAL EXHIBIT

Second Declaration of George T. Ligler

Exhibit 15

(Radiology, 2000;215:543-553.)
© RSNA, 2000

Computer Applications

Assessment of Visually Lossless Irreversible Image Compression: Comparison of Three Methods by Using an Image-Comparison Workstation¹

Richard M. Slone, MD, David H. Foos, MS, Bruce R. Whiting, PhD, Edward Muka, MSE, David A. Rubin, MD, Thomas K. Pilgram, PhD, Kevin S. Kohm, MSEE, Susan S. Young, MS, Paul Ho, MD and Dan D. Hendrickson, MS

¹ From the Mallinckrodt Institute of Radiology, Washington University School of Medicine, Barnes-Jewish Hospital, Box 8131, 510 S Kingshighway Blvd, St Louis, MO 63110. From the 1998 RSNA scientific assembly. Received December 30, 1998; revision requested February 23, 1999; revision received July 16; accepted August 18. Address correspondence to R.M.S. (e-mail: sloner@mir.wustl.edu).

► Abstract

PURPOSE: To determine the degree of irreversible image compression detectable in conservative viewing conditions.

MATERIALS AND METHODS: An image-comparison workstation, which alternately displayed two registered and magnified versions of an image, was used to study observer detection of image degradation introduced by irreversible compression. Five observers evaluated 20 16-bit posteroanterior digital chest radiographs compressed with Joint Photographic Experts Group (JPEG) or wavelet-based trellis-coded quantization (WTCQ) algorithms at compression ratios of 8:1–128:1 and x2 magnification by using (a) traditional two-alternative forced choice; (b) original-revealed two-alternative forced choice, in which the noncompressed image is identified to the observer; and (c) a resolution-metric method of matching test images to degraded reference images.

This Article

- [Abstract FREE](#)
- [Figures Only](#)
- [Full Text \(PDF\)](#)
- [Submit a response](#)
- [Alert me when this article is cited](#)
- [Alert me when eLetters are posted](#)
- [Alert me if a correction is posted](#)
- [Citation Map](#)

Services

- [Email this article to a friend](#)
- [Similar articles in this journal](#)
- [Similar articles in PubMed](#)
- [Alert me to new issues of the journal](#)
- [Download to citation manager](#)
- [Cited by other online articles](#)

PubMed

- [PubMed Citation](#)
- [Articles by Slone, R. M.](#)
- [Articles by Hendrickson, D. D.](#)

- ▲ [TOP](#)
- [Abstract](#)
- ▼ [Introduction](#)
- ▼ [MATERIALS AND METHODS](#)
- ▼ [RESULTS](#)
- ▼ [DISCUSSION](#)
- ▼ [References](#)

AX205595

RESULTS: The visually lossless threshold was between 8:1 and 16:1 for four observers. JPEG compression resulted in performance as good as that with WTCQ compression at these ratios. The original-revealed forced-choice method was faster and as sensitive as the two-alternative forced-choice method. The resolution-metric results were robust and provided information on performance above visually lossless levels.

CONCLUSION: The image-comparison workstation is a versatile tool for comparative assessment of image quality. At x2 magnification, images compressed with either JPEG or WTCQ algorithms were indistinguishable from unaltered original images for most observers at compression ratios between 8:1 and 16:1, indicating that 10:1 compression is acceptable for primary image interpretation.

Index terms: Computers, diagnostic aid • Data compression • Images, artifact, **.93, **.99² • Images, display, **.1215, **.99² • Images, processing, **.99² • Picture archiving and communication system (PACS)

► Introduction

For the past 2 decades, many researchers in radiology have predicted the replacement of hard copy-based, manually administered, diagnostic imaging operations by electronic picture archiving and communication systems (PACS). Re-engineering of radiology operations with PACS can dramatically improve health care delivery by enabling rapid distribution of images and information, improving resource utilization, and providing better service to caregivers (1,2). Although advances in technology have allowed successful demonstration of this concept, the high cost of systems (3) and the low comfort level of health care delivered with new and unfamiliar techniques has impeded widespread deployment. An important factor in the cost of PACS is the large amount of information contained in radiologic images, which results in terabytes of data that must be managed and distributed (1,2). The requirements for storage devices and networks thus constitute a substantial investment and ongoing costs to achieve the benefits of PACS.

Data compression, a technology that reduces the size of image files, provides immediate and substantial reduction in the cost of PACS deployment. Image compression techniques are designed to reduce data redundancy by means of special image coding and, as a result, can greatly reduce the effective amount of image data and, therefore, the volume of storage or transmission time required per image. Mathematically lossless compression techniques (compression and reconstruction with no loss of original data) result in compression factors on the order of 2:1 to 3:1 for radiologic images, which are insufficient to produce adequate reductions in transmission time or storage costs. To achieve these goals, compression factors on the order of 10:1 or higher are required, which implies that irreversible or lossy compression must be used; that is, some

▲ TOP
▲ Abstract
▪ Introduction
▼ MATERIALS AND METHODS
▼ RESULTS
▼ DISCUSSION
▼ References

information must be lost in the compression and reconstruction process. However, some loss of image data can be tolerated without affecting the visual interpretation of an image (4).

A major challenge in the adoption of lossy image compression in the medical community is to develop a body of research that supports the use of fewer data than are included in the full image for primary diagnostic interpretation. The characteristics of the human visual system are such that an image reconstructed after irreversible compression may appear indistinguishable from the original, and, thus, the compression is "visually lossless" (5). We believe that under these circumstances, the image is therefore diagnostically lossless; that is, image compression will have no effect on diagnostic interpretation. The purpose of our investigation was to compare three methods for the evaluation of compression artifacts by using a workstation designed to increase observer sensitivity to subtle differences, thereby arriving at a conservative and, we hope, widely accepted estimate of the visually lossless threshold. Image quality assessed by evaluating observers' perception of degradation as a function of compression ratio was the primary focus of our investigation.

► MATERIALS AND METHODS

Image Comparison Workstation

The image comparison workstation (ICW) was developed as a collaborative project between the Electronic Radiology Laboratory at our institution and the Health Imaging Research Laboratory of Eastman Kodak (Rochester, NY). The goal was to construct a workstation and software that would allow rapid processing and presentation of images on high-resolution (2,000 x 2,500-pixel) monitors. The ICW was designed specifically for the study of performance with lossy image compression techniques (6).

▲ TOP
▲ Abstract
▲ Introduction
• MATERIALS AND METHODS
▼ RESULTS
▼ DISCUSSION
▼ References

The workstation consisted of a personal computer (Kayak XU6/300; Hewlett Packard, Palo Alto, Calif) equipped with dual 300-MHz Pentium II processors (Intel, Santa Clara, Calif), a 9-Gbyte RAID (redundant array of inexpensive disks) disk array, and 512 Mbytes of random access memory), with model P1540 display cards (Metheus, Beaverton, Ore) driving a 21-inch-diagonal (53.3-cm-diagonal), 2,048 x 2,560-pixel, low-spatial-noise phosphor (P45), 71-Hz monitor (model DR 110; Data Ray, Westminster, Colo) with a maximum luminance of 220 candelas per square meter (luminance dynamic range of 650:1). The software application, which is implemented for the Windows NT operating system (Microsoft, Redmond, Wash), was developed by the Health Imaging Research Laboratory (Eastman Kodak). The user interface is shown in Figure 1.

AX205597

Figure 1. Direct screen capture shows the ICW graphical user interface. A small replication of the entire image is shown in

6/25/2006
5/25/2006



View larger version (94K):
[\[in this window\]](#)
[\[in a new window\]](#)

the lower left-hand corner, where the white box defines the portion of the image displayed in the top 80% (2,048 x 2,048 pixels) of the monitor. The set of resolution images for comparison is listed in the box on the right for test case 18-5. The "Prev" and "Next" buttons in the "Test Image" area change the test image being compared. The "Prev" and "Next" buttons in the "Control Image" area change the resolution reference image. This can also be changed by using the wheel on the mouse or the up and down arrows in the "Best Match" area. The "Flicker" area offers selection of "Auto" or "Manual" for automatic or manual control of flicker, respectively, and a slider to set the flicker rate in the automatic control mode. The "Zoom In" area offers a choice of x1, x2, or x4 magnification. The x2 magnification is displayed in this example. The "Mark as Best" button records the control image selected as the best match to the test case. A warning is displayed if the user tries to move to the next case without recording a choice.

The approach used in this study was to compare two versions of an image on a single monitor by using an interactive soft-copy feature. Inherent in the design was the use of "flicker," which is defined as sequential display of two registered images on the same monitor. This method was used to exploit the observer's temporal sensitivity to differences in the image, because the human visual system is naturally drawn to changes in structure or brightness. This technique allows detection of subtle differences and provides a mechanism for comparing image quality loss caused by different kinds of distortion. The observer has direct control of flicker and can set it to automatically change images at up to five times per second; alternatively, the user can use a manual mode, which allows the observer to selectively toggle between the two images, as desired.

The current software allowed x1 magnification to simulate clinical application, while x2 and x4 magnifications were available to help improve detection of subtle differences. A small representation of the entire image with the area chosen for magnification was displayed in the lower left-hand corner of the monitor. The portion of the image displayed for evaluation occupied the upper 80% of the screen, as shown in [Figure 1](#), and measured 30 x 30 cm. The observer could change the region of the image displayed by panning with the mouse on the image representation in the lower left-hand corner. Thus when x2 or x4 magnification was used, the observer was free to study any desired segment of the full image. A wide assortment of information could be recorded automatically by the computer as an observer worked through an experiment to compare a series of images and to make choices.

Images

The image set used in this study consisted of 20 digital posteroanterior chest radiographs obtained from the outpatient admitting area at our institution. A commercial selenium detector system (Thoravision; Philips Medical Systems, Shelton, Conn) was used to obtain the images. Images normally have an addressable area of 2,048 x 2,560 pixels with a pixel size of 0.2 mm. The use of

digital chest images in the context of primary interpretation is well supported by research results (7–9) on radiologists' preference and performance as assessed with receiver operating characteristic analysis for comparison with state-of-the-art, wide-latitude, dual screen-film images. The images selected for this investigation included studies in men and women with pneumonia, pulmonary nodules, interstitial lung disease, mediastinal masses, catheters, or implanted hardware.

Image Compression

There are many compression techniques available and much research on algorithm improvement (10,11); however, interoperability is crucial to a successful PACS, so we have focused our efforts on existing standards defined by the Joint Photographic Experts Group (JPEG) (12). JPEG baseline is the most widely available block discrete cosine transform algorithm. Advantages include wide availability, interoperability with other JPEG-compliant encoding and decoding software, reasonably fast "run times," and widespread vendor support. It is the only compression algorithm sufficiently documented to be proposed by the National Electrical Manufacturers Association (13) as a standard for Digital Imaging and Communications in Medicine, or DICOM.

Other important methods to investigate are wavelet-compression techniques, because these techniques have special features and functionality (14–16). Driven by broad interest, the JPEG 2000 Committee was established to formulate a new standard, ostensibly to be based on wavelet compression. We thus included the wavelet-based trellis-coded quantization (WTCQ) algorithm developed at the University of Arizona as a representative example of this class of algorithm (17).

The original radiograph (postprocessed, relative log luminance data) was retrieved from an optical disk and linearly transformed from a 0–30,000 scale to a 0–4,095 scale to match the 12-bit requirements of the compression algorithms. The image compression rates were selected on the basis of pilot data and corresponded to 2.00, 1.50, 1.00, 0.75, 0.50, 0.25, and 0.125 bits per pixel. Because the original image was created with 2 bytes per pixel, as is typical for most commercial digital radiographic systems wherein "byte-packing" is not used, we calculated, for the convenience of the reader, a compression ratio defined as 16 bits divided by the number of compressed bits per pixel. Representative examples are shown in [Figure 2](#).




Figure 2. Compression artifacts. Compressed (JPEG, left column; WTCQ, right column) and noncompressed digital (middle column) posteroanterior radiographs of a selected region of interest show the effects of compression at ratios of 8:1 (top row), 16:1 (middle row), and 128:1 (bottom row). With both JPEG and WTCQ algorithms, the image compressed at 8:1 is indistinguishable from the original. At a compression ratio of 128:1, the manifestation of "tiling" or "blocking" artifacts on the JPEG image and blurring on the WTCQ image are readily apparent.

[View larger version](#)

AX205599

(105K):
[\[in this window\]](#)
[\[in a new window\]](#)

Observers

The five independent observers included two imaging scientists (B.R.W., E.M.) with extensive experience in image processing and display and three board-certified radiologists, including specialists in chest (R.M.S.), musculoskeletal (D.A.R.), and general (P.H.) radiology. The introductory training session for each observer included a discussion of the purpose and objectives of the evaluation, a description of the protocol, and an online walk-through of the evaluation procedure and operation of the ICW.

Image Evaluation Methods

Three methods were chosen for image comparison and evaluation. All three were conducted by using the ICW at x2 magnification, which allowed 25% of the total image area to be viewed at a time. Although it is possible to pan (roam) and view the entire image, observers limited their observations to the upper right-hand quadrant of each image for this study.

Twenty test "folders" (computer directories) were prepared and randomized for each reader. Each folder used a different image and contained 16 randomized image replicates, including two unaltered original radiographs and images compressed with each of the seven ratios for the two compression algorithms. The order in which the five observers performed the three experiments was random. The amount of time needed for each reading session, reader confidence for each decision, and representative reading distances were manually recorded.

Two-alternative forced choice.—The 16 test images in each folder were paired with a control image (unaltered original) and were presented sequentially without identification (one pair at a time), with the observer toggling between them in a rapid fashion. The observer was asked to choose the image with "better quality" and was forced to choose even if the observer perceived no difference between the images. The observer was asked to record both the image selected and the decision confidence by using a three point scale: score of 0, uncertain or guessing; score of 1, confident; score of 2, very confident. No reader feedback was provided. With this method, a visually lossless level was indicated when responses were evenly divided between test and control images or matched the score distribution for the original-original image pair.

Original-revealed forced choice.—The same 20 test folders were used to conduct a modified two-alternative forced-choice experiment. As in the traditional forced-choice method, each test image was paired with an unaltered original and was presented sequentially (one pair at a time) on the ICW, with the observer toggling between them. The difference between this task and two-alternative forced choice was that the original image was identified for the reader. The observer was asked to compare the test image to the original and decide if the images were equivalent or if there was visible artifact, loss of fidelity, or degradation in the quality of the test image. The

observer knew that sometimes two originals would be shown. We refer to this method as the original-revealed forced-choice method. In an investigation of the human visual system, Gur et al (18) advocate comparison of test images with a known original image to maximize sensitivity to distortion. With this method, a visually lossless level was indicated when the percentage of test images rated as equivalent approximated 100% or matched the score distribution for the original-original image pairs.

Spatial resolution metric.—In this experiment, a test image that had been prepared by applying the compression algorithm was compared to a set of reference images that had been prepared by degrading the image in a controlled manner. A set of 15 reference images was prepared for each test image by using "blur" as the metric. Blur was introduced by degrading the spatial resolution of the image while maintaining image size. This bandwidth reduction was performed by transforming the image into the frequency domain, applying a set of 14 power-law filters, then "back-transforming" the filtered data into the spatial domain. Power-law filters were selected to provide rectangular bandpass characteristics in the frequency domain while producing minimal distortion in the spatial domain (6).

Because the image is two-dimensional while the filters are separable (one-dimensional), the square of the bandwidth reduction can be used as a measure of two-dimensional spatial resolution. Figure 3 shows the full range of grades of spatial filtration applied to a representative image. For example, a bandwidth reduction factor of 1.25 (grade 4), which corresponded to a 20% reduction in bandwidth, can be thought of as displaying 80% of the original pixels, which were then magnified to the full display size, resulting in a spatial resolution of 64%. This would be analogous to displaying a 2,048 x 2,048-pixel image on a 1,600 x 1,600-pixel monitor. Similarly, a bandwidth reduction factor of 2 (grade 7) corresponded to displaying half the pixels, or 25% resolution, and was analogous to displaying a 2,048 x 2,048-pixel image on a 1,024 x 1,024-pixel monitor.




Figure 3. Resolution metric. Composite image shows the spatial-resolution scale applied to a representative portion of a radiograph. Reduction in resolution progresses from grade 1 (top left and bottom right: unaltered original, 100%) to grade 15 (bottom row, second from right: 4%). The percentages are a measure of two-dimensional spatial resolution, which was determined on the basis of the square of the bandwidth reduction implemented with power-law filters and can be thought of as the percentage of pixels displayed.

[View larger version](#)

(100K):

[\[in this window\]](#)

[\[in a new window\]](#)

AX205601

The test and reference images were presented sequentially for comparison at the ICW, with the observer toggling between them in a rapid fashion. The observer changed the comparison image by using the wheel on the mouse to select from the set of 15 reference images. The observer selected the reference image that most closely matched the test image in terms of clinical utility.

Data Analysis

Data took the form of reader decisions: For two-alternative forced-choice experiments, observers decided which of the two images appeared to be superior; for original-revealed forced-choice experiments, observers decided if an image was equivalent to the original or degraded; for resolution-metric tests, observers decided which level of blurring most closely matched, with respect to clinical utility, the level of compression. Analysis took the form of calculation of the proportion of decisions in a given category, with 95% confidence limits (19). Time needed to complete image comparison and reading distance were evaluated by comparing means.

► RESULTS

When making comparisons, observers reported that their focus of attention included structural detail, particularly bone edges and trabecular patterns, and areas of uniform opacity such as the soft tissues of the chest wall. Observers differed in their ability to detect degraded images, but when results from all observers were combined, a fairly clear pattern was found. The mean results for two-alternative forced-choice, original-revealed forced-choice, and spatial resolution-metric experiments for all cases and observers are shown in [Table 1](#). Individual results are shown in [Table 2](#).

▲ TOP
▲ Abstract
▲ Introduction
▲ MATERIALS AND METHODS
• RESULTS
▼ DISCUSSION
▼ References

View this table: TABLE 1. Combined Results for All Observers and Methods

[\[in this window\]](#)

[\[in a new window\]](#)

View this table: TABLE 2. Visually Lossless Threshold Estimates for Individual
[\[in this window\]](#) Observers

[\[in a new window\]](#)

AX205602

Two-Alternative Forced Choice

The overall results for the two-alternative forced-choice experiments are presented in [Figure 4](#). When presented with images that were indistinguishable, observers guessed, which resulted in a chance, or approximately 50:50, distribution. As part of the two-alternative forced-choice

experiment, observers were presented with 40 pairs of original images, one denoted as the "test" image and one as the "control" image. The overall response rate for selecting the test or control image as the better image (when in fact both were identical) was 47% or 53%, respectively.




Figure 4. Bar graph shows the combined results for five observers in the two-alternative forced-choice experiments, with the percentage of test images judged to be superior to the original image. The horizontal line at 50% represents the expected result for a purely random selection. Error bars = 95% CIs, gray bars = WTCQ images, white bars = JPEG images.

View larger version (29K):

[\[in this window\]](#)

[\[in a new window\]](#)

AX205603

The mean responses for images compressed at 8:1 and 11:1 with both algorithms fall within the 95% CIs of the responses for the original images. The 8:1 compressed images were selected as the better image slightly more often than the original for both JPEG (57%) and WTCQ (65%) images. The responses for both JPEG and WTCQ images compressed at 11:1 indicate a slight tendency for observers to select the original image as the better image. The responses for images at 16:1 compression indicate that observers differentiated WTCQ images from original images more frequently than they differentiated JPEG images and that the visually lossless compression threshold was crossed for both. At compression ratios of 21:1 and higher, observers consistently (more than 95% of the time) chose the original as the better image, which indicates that the compression artifact was clearly evident.

These results are supported by the mean confidence scores for two-alternative forced-choice experiments (Table 1). Observer confidence was low (mean score of 0.2) when the choice was between images in an original-original pair and remained low with 8:1 and 11:1 compressed images. Observer confidence increased (mean score > 1.0) when evaluating images compressed at 16:1 and became substantially higher (mean score > 1.7) at compression ratios of 21:1 and higher.

Individual observers trends varied, as shown in Table 2. With JPEG images, the 8:1 compression images were indistinguishable from the original images for four observers, and the 11:1 compression images were indistinguishable for three observers. Observer C, who also had the closest mean viewing distance (as close as 8 cm), noted degradation in all compressed images except for a few compressed at 8:1. Although observer C noted no structural degradation, he noted a change in some individual pixels. With WTCQ compression, the pattern for all five observers was similar and suggested the presence of a visually lossless threshold at x2 magnification for

observers.

Original-revealed Forced Choice

In this experimental design, the observer was presented with an original image and was asked if the test image was indistinguishable or if there was any visible artifact, degradation, or loss of fidelity. Observers were told that there were several unaltered original images included with the test images. When presented with a pair of original images, the observer declared the images to be indistinguishable 95% of the time. The composite results for all observers are shown in [Figure 5](#). As with the two-alternative forced-choice results, the greatest change in performance was found between 11:1 and 16:1 compression for both JPEG and WTCQ algorithms.




Figure 5. Bar graph shows the combined results for five observers in the original-revealed forced-choice experiments, with the percentage of test images classified as equivalent to an original image. Error bars = 95% CIs, gray bars = WTCQ images, white bars = JPEG images.

[View larger version \(31K\):](#)

[\[in this window\]](#)

[\[in a new window\]](#)

Review of the individual results showed a tendency for two observers to detect degradation in about half of the WTCQ images compressed at 11:1 and a clear trend for four observers to detect degradation in such images compressed at a ratio higher than 16:1. For JPEG images, a larger percentage of images were considered to be degraded at a given compression ratio overall, but this was due primarily to the results of observer C, who detected a difference in all compressed images except for a few compressed at 8:1. Of the other four observers, one detected degradation in about half the JPEG images compressed at 11:1, but none detected degradation in images compressed at 8:1. The mean percentages for these four observers were 94% for JPEG images compressed at 8:1 and 90% for images compressed at 11:1—in both cases, close to the 95% rate found for original-original pairs. At compression ratios of 21:1 and higher with either JPEG or WTCQ, the degradation was detected over half the time by four of the observers.

Spatial Resolution Metric

Fifteen reference images with controlled degradation were available for comparison with each test image, including grade 1 degradation (ie, an unaltered original image). Although observers were able to score intermediate grades, this seldom occurred, which suggests that the existing choices offered were sufficient. Observers reported the ability to consistently select a reference image that matched the degradation in the compressed image, although this was easier with WTCQ images

than with JPEG images at high compression ratios.


The original images were matched with grade 1 images (unaltered original) on the resolution scale 96% of the time; with grade 2, 3% of the time; and with grade 3, 1% of the time. At a compression ratio of 8:1, JPEG images were matched with grade 1 images 94% of the time; with grade 2 images, 5% of the time; and with grade 3, 1% of the time. At a ratio of 11:1, JPEG images were matched with grade 1 images 71% of the time; with grade 2 images, 15% of the time; with grade 3 images, 11% of the time; and with grade 4–6 images, 1% of the time each. At a compression ratio of 8:1, WTCQ images were matched with grade 1 images 96% of the time; with grade 2 images, 3% of the time; and with grade 4 images, 1% of the time. At a compression ratio of 11:1, WTCQ images were matched with grade 1 images 67% of the time; with grade 2 images, 8% of the time; with grade 3 images, 13% of the time; and with grades 4–6 images, 3% of the time each.

The mean percentage resolution matched with each level of compression for the two algorithms is shown in Table 1. There was a progressive decrease in resolution as the compression level increased. As with the two-alternative forced-choice and original-revealed forced-choice experiments, the greatest initial change for both algorithms occurred between compression ratios of 11:1 and 16:1. Although the data were not highly precise, at low compression levels JPEG images were matched with a higher resolution than were WTCQ images, whereas the opposite occurred for the highest compression ratios.


The relationship between compression and resolution was further assessed by plotting the data for JPEG images (Fig 6a) and WTCQ images (Fig 6b). Transformation of the data onto a log scale provided a distribution suitable for linear regression. The log of the compression ratio was plotted on the x axis. In this format, 1.0 corresponded to a compression ratio of 10:1; 1.2, to a compression ratio of 16:1; and so forth. The log of the reciprocal of the percentage spatial resolution was plotted on the y axis and ranged from 0 for the unaltered original to 1.4 for the bandwidth reduction factor of 5.0 (4% resolution). Linear regression was used to calculate the line of best fit. The x intercept was a predictor of the visually lossless threshold. This method was used to calculate the expected visually lossless threshold for each observer by using data for compression ratios higher than the visually lossless level, namely, 16:1 and higher. The calculated threshold and r^2 value for the linear regression fit to the data are shown in Table 2 for each reader and both algorithms.

AX205605

Figure 6a. Resolution-metric matching with (a) JPEG images and (b) WTCQ images. Scatterplots show the relationship between compression ratio (CR) and percentage resolution for observer A. These data were obtained by matching the blur image set with 20 test images at each of five compression ratios (16:1, 21:1, 32:1, 64:1, and 128:1). The area of each dot is proportional to the number of superimposed data points. The straight lines are the regression lines. (a) The x



View larger version (29K):
[\[in this window\]](#)
[\[in a new window\]](#)



View larger version (28K):
[\[in this window\]](#)
[\[in a new window\]](#)

intercept for JPEG images, determined by means of linear regression on the 100 data points, indicates that the estimated visually lossless threshold is 13:1. (b) The x-intercept for WTCQ images, determined by means of linear regression on the 100 data points, indicates that the visually lossless threshold is 11:1.

Figure 6b. Resolution-metric matching with (a) JPEG images and (b) WTCQ images. Scatterplots show the relationship between compression ratio (CR) and percentage resolution for observer A. These data were obtained by matching the blur image set with 20 test images at each of five compression ratios (16:1, 21:1, 32:1, 64:1, and 128:1). The area of each dot is proportional to the number of superimposed data points. The straight lines are the regression lines. (a) The x intercept for JPEG images, determined by means of linear regression on the 100 data points, indicates that the estimated visually lossless threshold is 13:1. (b) The x-intercept for WTCQ images, determined by means of linear regression on the 100 data points, indicates that the visually lossless threshold is 11:1.

Linear regression also was performed for each observer with systematic elimination of one compression level and again with only the 16:1, 32:1, and 64:1 data. These data manipulations had almost no effect on the x intercept. The mean visually lossless compression ratio was 12:1 for JPEG images and 11:1 for WTCQ images, in all cases.

Viewing Distance and Reading Time

Viewing distance and reading time varied among observers (Table 2). The resolution-metric task required the greatest time commitment, taking about twice as long to complete as the two-alternative forced-choice task. The original-revealed forced-choice task was the quickest method, taking, on average, 25% less time than the two-alternative forced-choice task.

Four observers wore prescription eyeglasses, and one (observer D) wore nonprescription reading glasses. Observer C positioned himself much closer to the images than did the other four observers. This observer was very nearsighted and, with prescription lenses, was able to focus

much closer than the typical viewer. We have already noted the differences in some of the results of observer C.

► DISCUSSION

Reversible, or mathematically lossless, image compression provides an inadequate reduction in the amount of data to provide substantial engineering advantages or cost reductions for image transmission or storage. Irreversible, or lossy, image compression is needed to achieve these goals. In certain circumstances, visually lossy images may be diagnostically lossless; although image compression artifacts are detectable, their presence does not affect diagnostic performance (20). Several studies in which receiver operating characteristic analysis was used (21) have shown this to be true (22–24).

▲ TOP
▲ Abstract
▲ Introduction
▲ MATERIALS AND METHODS
▲ RESULTS
▪ DISCUSSION
▼ References

Although diagnostically lossless criteria would likely allow relatively high degrees of compression and substantial cost savings, the presence of perceivable artifact reduces acceptance among skeptical radiologists, and the validity of diagnoses based on lossy images in medicolegal proceedings is yet to be determined. In addition, receiver operating characteristic analyses are time-consuming and expensive and can usually be used to address only narrowly defined tasks. Multiple receiver operating characteristic studies would be required to provide results that cover a broad range of potential abnormalities with confidence. In light of this, we chose to concentrate on a more conservative criterion—namely, visually lossless compression for the full range of texture and density gradients in the image—in the belief that compressed images that are indistinguishable in any way from an original image are diagnostically lossless and would be readily acceptable even by skeptical radiologists.

The ICW proved to be a powerful tool for conducting observer studies. A large number of image sets could be evaluated efficiently. The alternating presentation of registered images at the ICW was intended to maximize viewer sensitivity to subtle image compression artifacts. Similarly, viewing distance was unconstrained, and images were magnified by a factor of two. Both of these latter factors should increase the conspicuity of subtle detail not visible when the region of interest occupied a smaller viewing angle. Results of previous studies (25,26) have shown decreases in detection of degradation due to image compression as viewing distance increases.

Display of images magnified by a factor of two also has clinical relevance, because most PACS stations have a "magnifying glass" tool. Radiologists may use this tool to evaluate isolated areas of the image at x2 magnification to detect subtle disease, particularly pneumothorax, fracture, and interstitial lung disease. We believe, therefore, that the use of magnification, close viewing distance, and flicker to exploit an observer's temporal sensitivity between image differences should result in a conservative and, we hope, widely accepted estimate of the visually lossless threshold.

We compared three methods for studying observer detection of image degradation. The two-alternative forced-choice and original-revealed forced-choice methods were similar. The difference was subtle but important. In the former method, when confronted with two original images or a visually lossless image, the observer was forced to guess, thus selecting the image that was compressed but indistinguishable from the original image 50% of the time. As the degradation became more apparent, selection of the compressed image decreased toward 0%. In the original-revealed forced-choice method, the observer would be expected to declare two original or visually lossless images to be equivalent 100% of the time. However, sometimes an original was judged to be degraded because the observer was intent on detecting the subtlest differences and may have "over read" the image. This was observed in our original-revealed forced-choice results, where observers declared the unaltered test image to be degraded 5% of the time and were thus operating at a low threshold for reporting degradation.

With the original-revealed forced-choice method, as the degradation became more apparent, judgment of the compressed image as equivalent to the original decreased toward 0%. Thus, the expected results ranged from 0% to 50% with the two-alternative forced-choice method and from 0% to 100% with the original-revealed forced-choice method. This has certain advantages for the display and analysis of the data. The patterns shown in the data suggest that observers detected degradation in the compressed images as frequently when using the original-revealed forced-choice method as they did with the two-alternative forced-choice method, which implies that the former technique is as sensitive as the latter. The original-revealed forced-choice experiment was conducted in less time by the observers and thus may have provided results more efficiently.

In contrast, an advantage of the two-alternative forced-choice method is that it can reveal trends in preference rather than just provide information about detection of differences. Of particular interest are the results for comparisons with images compressed at 8:1. For both JPEG and WTCQ, the 8:1 images were considered to be better than the original image 55%–65% of the time. Although this was a slight deviation from the expected chance result of 50% for visually lossless images, the trend was the opposite of that with all other levels of compression, where the compressed image was judged to be better less than 50% of the time. We hypothesize that low compression levels have the same effect as a low-pass filter, because the image is smoothed and the conspicuity of image noise is effectively decreased. Thus, the original-revealed forced-choice method was sensitive for detecting differences, but only with the two-alternative forced-choice method did the results reflect a preference for a test image over an original image.

Despite expectations for improved performance with wavelet-based algorithms, we found that the JPEG baseline algorithm resulted in performance that was as good as, if not better than, performance with the WTCQ algorithm implemented at low compression ratios. This is critical, because JPEG is a current standard that permits interoperability in a PACS environment. The data did suggest that WTCQ was better at very high compression ratios. These results are probably related to the way these two algorithms handle the data and the manner in which the artifact is manifested (ie, primarily as "tiling" or "blocking" with JPEG and as blurring with WTCQ).

Although radiographs of different body parts are likely to emphasize different compression artifacts, the chest radiograph demonstrates a broad range of structural and tonal characteristics that provide the opportunity to note degradation in soft tissues of uniform opacity and intricate trabecular and pulmonary parenchymal detail. Thus, although our study was limited to posteroanterior chest radiographs, we believe the results should be representative of other projection radiographs.

The resolution-metric method was intended to project the image characteristics of the compression algorithm onto a quantifiable dimension, thus providing a basis for comparison of various compression techniques. Observers reported being able to consistently select a reference image that matched the degradation in the compressed image, although at high compression ratios this was easier with WTCQ images than with JPEG images. We believe this is because the artifact introduced by the WTCQ algorithm is similar to the degradation introduced by blurring. At high compression ratios, it was more difficult to equate the tiling or blocking artifacts introduced by JPEG with blur artifacts introduced by WTCQ.

The use of spatial blurring by means of bandwidth limitation in the frequency domain proved to be a reasonable choice for a matching metric. Further exploration of degradation mechanisms for reference images might include different frequency filtering, variation of quantization levels, addition of random noise (both white noise and quantum mottle), or a combined blur-noise operation that follows a relationship similar to that of conventional x-ray detectors (ie, a speed-sharpness trade-off similar to that of screen-film or storage phosphor images).

This technique required more time than did the forced-choice methods but provided more information, particularly about the relative "performance" of algorithms at compression ratios above the visually lossless threshold. Such a comparison would be critical in a comparison of a new algorithm with an accepted one in the visually lossy range. In addition, we found that data from the resolution matching method could be used to estimate the visually lossless threshold by using linear regression. The results were nearly identical to those obtained with the two-alternative forced-choice and original-revealed forced-choice methods, which showed that the visually lossless threshold for the conservative viewing situation (superimposed images, x2 magnification, close viewing conditions) was greater than 10:1 for both algorithms, with JPEG images resulting in slightly better performance than WTCQ images. We also found this technique to be robust, with the number of data points and specific compression levels used having little effect on the predicted value.

The ICW is a versatile and powerful tool for comparative assessment of image quality. Sequential registered display of magnified images should help optimize observer sensitivity to differences and improve detection of subtle degradation, resulting in conservative estimates. The ICW monitor is a component of currently available commercial PACS systems, which makes translation of the results to a true clinical environment practical. These studies were conducted in the context of primary, rather than secondary, interpretation, and the methods were robust in that

they could be used with various equipment and acquisition, presentation, display, and viewing tasks.

The objective of our research was to establish a basis for acceptance of irreversible image compression for primary interpretation of diagnostic images. We believe that diagnostic loss can be avoided by using visually lossless levels of compression and that these compression levels are high enough to provide important time and cost savings and thus serve as a critical component for improved health care delivery in the PACS environment. Our results suggest that the JPEG baseline algorithm results in performance that is as good as that which results from a more complex wavelet-compression algorithm and that 10:1 image compression is visually lossless for most observers and is, therefore, acceptable for primary image interpretation without risk of affecting diagnosis. The composite results shown in Figures 4 and 5 demonstrate that most observers consistently detect image degradation at compression ratios of 21:1 and higher. This is supported by other researchers (27) who have suggested that 10:1 compression does not influence detection of subtle interstitial abnormalities but that important information may be lost at a ratio of 20:1, particularly when the images are interpreted by experienced thoracic radiologists.

Our research design, which used magnification, unrestricted viewing distance, and superimposition of registered images, was intended to produce a conservative estimate of acceptable image compression levels. Our hope is that comfort with visually lossless but mathematically degraded data sets will open the path for studies in a routine clinical environment, where less conservative constraints would support the use of more aggressive compression with its accompanying benefits. Future work is needed to address less stringent visually lossy but diagnostically lossless levels of compression by assessing diagnostic performance outcomes.

► Acknowledgments

We are grateful to G. James Blaine, DSc, Jerome Cox, PhD, and R. Gilbert Jost, MD, at the Mallinckrodt Institute of Radiology, Washington University School of Medicine (St Louis, Mo), for their critical insight, thoughtful suggestions, and support in facilitating this collaborative investigation.

► Footnotes

²
**. Multiple body systems ■

AX205610

Abbreviations: ICW = image compression workstation JPEG = Joint Photographic Experts Group
PACS = picture archiving and communication system WTCQ = wavelet-based trellis-coded quantization

Author contributions: Guarantor of integrity of entire study, R.M.S.; study concepts, B.R.W.,

E.M.; study design, E.M., D.H.F., S.S.Y., K.S.K.; definition of intellectual content, E.M., B.R.W., D.H.F.; literature research, R.M.S.; experimental studies, D.H.F., K.S.K., S.S.Y.; data acquisition, R.M.S., E.M., P.H., D.A.R., B.R.W.; data analysis, R.M.S., T.K.P.; statistical analysis, T.K.P.; manuscript preparation, R.M.S.; manuscript editing, D.D.H., R.M.S., B.R.W., E.M.; manuscript review, D.A.R., D.H.F.


► References

1. Blaine GJ, Jost RG, Martin L, Weiss DA. Information and image integration: Project Spectrum. Proc SPIE 1998; 3339:430-439.
2. Kahn MG. Enterprise-wide clinical data integration. In: Brennan P, Schneider SJ, Tornquist E, eds. Information networks for community health. New York, NY: Springer-Verlag, 1997; 41-45.
3. Beard D, Parrish D, Stevenson D. A cost analysis of film image management and four PACS based on different network protocols. J Digit Imaging 1990; 3:108-118. [\[Medline\]](#)
4. Cox GG, Cook LT, Insana MF, et al. The effects of lossy compression on the detection of subtle pulmonary nodules. Med Phys 1996; 23:127-132. [\[Medline\]](#)
5. Daly S. Application of a noise-adaptive contrast sensitivity function to image data compression. Opt Eng 1990; 29:977-987.
6. Foos DH, Slone RM, Whiting B, et al. Dynamic viewing protocols for diagnostic image comparison. Proc SPIE 1999; 3663:108-120.
7. Woodard PK, Slone RM, Gierada DS, Reiker GR, Pilgram TK, Jost RG. Chest radiography: depiction of normal anatomy and pathologic structures with selenium-based digital radiography versus screen-film radiography. Radiology 1997; 203:197-201. [\[Abstract\]](#)
8. Woodard PK, Slone RM, Sagel SS, et al. Detection of CT-proved pulmonary nodules: comparison of selenium-based digital and conventional screen-film chest radiographs. Radiology 1998; 209:705-709. [\[Abstract\]](#)
9. Floyd CE, Baker JA, Chatos HG, Delong DM, Ravin CE. Selenium-based digital radiography of the chest: radiologists preference compared with film-screen radiographs. AJR Am J Roentgenol 1995; 165:1353-1358. [\[Abstract\]](#)
10. Jones PW, Daly S, Gaborski RS, Rabbani M. Comparative study of wavelet and DC decompositions with equivalent quantizations and encoding strategies for medical images. Proc SPIE 1995; 2431:571-582.
11. Eckert MP. Lossy compression using wavelets, block DCT, and lapped orthogonal transforms optimized with a perceptual model. Proc SPIE 1997; 3031:339-350.
12. Pennebaker WB, Mitchell JL. JPEG still image data compression standard New York, NY: Van Nostrand Reinhold, 1993.
13. Digital Imaging and Communications in Medicine (DICOM). V. Data structures and encoding. NEMA standards publication no. PS 3.5-1996 Rosland, Va: National Electrical Manufacturers Association, 1996.
14. Saipetch P, Ho B, Ma M, Chuang K, Wei J. Radiological image compression using wavelet transform with arithmetic coding. Proc SPIE 1994; 2164:449-459.
15. DeVore R, Jawerth B, Lucier B. Image compression through wavelet transform coding. IEEE Trans Info Theory 1992; 38:719-746.
16. Goldberg M, Pivovarov M, Mayo-Smith WW, et al. Application of wavelet compression to digitized radiographs. AJR Am J Roentgenol 1994; 163:463-468. [\[Abstract\]](#)

▲ TOP
▲ Abstract
▲ Introduction
▲ MATERIALS AND METHODS
▲ RESULTS
▲ DISCUSSION
• References

17. Sriram P, Marcellin MW. Image coding using wavelet transforms and entropy-constrained trellis coded quantization. *IEEE Trans Image Proc* 1995; 4:725-733.
18. Gur D, Rubin D, Kart B, et al. Forced choice and ordinal discrete rating assessment of image quality: a comparison. *J Digit Imaging* 1997; 10:103-107.
19. Fleiss JL. Statistical methods for rates and proportions New York, NY: Wiley, 1981.
20. Mori T, Nakata H. Irreversible data compression in chest imaging using computed radiography: an evaluation. *J Thorac Imaging* 1994; 9:23-30. [\[Medline\]](#)
21. Metz CE. ROC methodology in radiologic imaging. *Invest Radiol* 1986; 21:720-733. [\[Medline\]](#)
22. MacMahon H, Doi K, Sanada S, et al. Data compression: effect on diagnostic accuracy in digital chest radiography. *Radiology* 1991; 178:175-179. [\[Abstract\]](#)
23. Aberle DR, Gleeson F, Sayre JW, et al. The effect of irreversible image compression on diagnostic accuracy in thoracic imaging. *Invest Radiol* 1993; 28:398-403. [\[Medline\]](#)
24. Savchenko V, Erickson BJ, Palisson PM, et al. Detection of subtle abnormalities on chest radiographs after irreversible compression. *Radiology* 1998; 206:609-616. [\[Abstract\]](#)
25. Cox JE, Muka E, Wang X, Blaine GJ. Factors affecting the selection of compression algorithms for projection radiography. *Proc SPIE* 1997; 3031:256-264.
26. Pilgram TK, Slone RM, Muka E, Cox JR, Blaine GJ. Perceived fidelity of compressed and reconstructed radiological images: a preliminary exploration of compression, luminance, and viewing distance. *J Digit Imaging* 1998; 11:168-175. [\[Medline\]](#)
27. Kido S, Ikezoe J, Kondoh H, et al. Detection of subtle interstitial abnormalities of the lungs on digitized chest radiographs: acceptable data compression ratios. *AJR Am J Roentgenol* 1996; 167:111-115. [\[Abstract\]](#)

This article has been cited by other articles:



Radiology


[HOME](#)

J. P. Ko, H. Rusinek, D. P. Naidich, G. McGuinness, A. N. Rubinowitz, B. S. Leitman, and J. M. Martino

Wavelet Compression of Low-Dose Chest CT Data: Effect on Lung Nodule Detection

Radiology, July 1, 2003; 228(1): 70 - 75.

[\[Abstract\]](#) [\[Full Text\]](#) [\[PDF\]](#)



Radiology

[HOME](#)

CT Colonography with Teleradiology: Effect of Lossy Wavelet Compression on Polyp Detection—Initial Observations

Radiology, August 1, 2001; 220(2): 387 - 392.

[\[Abstract\]](#) [\[Full Text\]](#) [\[PDF\]](#)

AX205612

Radiographics

[HOME](#)

R. H. Wiggins III, H. C. Davidson, H. R. Harnsberger, J. R. Lauman, and P. A. Goede

Image File Formats: Past, Present, and Future



RadioGraphics, May 1, 2001; 21(3): 789 - 798.
[\[Abstract\]](#) [\[Full Text\]](#)

This Article

- ▶ [Abstract FREE](#)
- ▶ [Figures Only](#)
- ▶ [Full Text \(PDF\)](#)
- ▶ [Submit a response](#)
- ▶ [Alert me when this article is cited](#)
- ▶ [Alert me when eLetters are posted](#)
- ▶ [Alert me if a correction is posted](#)
- ▶ [Citation Map](#)

Services

- ▶ [Email this article to a friend](#)
- ▶ [Similar articles in this journal](#)
- ▶ [Similar articles in PubMed](#)
- ▶ [Alert me to new issues of the journal](#)
- ▶ [Download to citation manager](#)

PubMed

- ▶ [PubMed Citation](#)
- ▶ [Articles by Stone, R. M.](#)
- ▶ [Articles by Hendrickson, D. D.](#)

HOME HELP FEEDBACK SUBSCRIPTIONS ARCHIVE SEARCH TABLE OF CONTENTS
RADIOLOGY RADIOGRAPHICS RSNA JOURNALS ONLINE

AX205613